Should You Believe the Truth?*

Jennifer Rose Carr

Abstract

It's often treated as a truism that we objectively epistemically ought to believe, or have maximal credence in, the truth. This claim is open to several interpretations. I explore a variety of *prima facie* plausible deontic semantics for the claim, and argue that each generates unacceptable verdicts. (The objections I raise are independent of human cognitive limitations.) I conclude that there's no genuine Truth Norm on belief: the objective 'ought' doesn't extend to epistemology. I extend the argument to weaker putative objective epistemic norms, according to which one ought to hold beliefs or credences that approximate the truth as closely as possible. While we can assign alethic value to beliefs or possibilities, we can't generate objective norms on this basis. This has significant ramifications in other areas of epistemology; I discuss the example of epistemic decision theory.

Epistemic rationality seems like it must have something to do with truth. Epistemic reasons to believe that p are considerations that provide evidence that p is true. What makes a kind of evaluation *epistemic* is somehow intimately tied with what's true or likely true: even if believing that p will make you happier or healthier, that doesn't make it epistemically rational for you to believe that p, because those considerations are independent of truth.

Working toward formulating a rigorous characterization of the connection between rationality and truth has been one of the core projects in epistemology for a very, very long time. Many epistemological theories are committed one or more of the following theses:

1. Truth is the aim of belief.

2. A belief is correct (in a normative sense) just in case it's true.

*Note: this is not the final draft! Thanks to Dan Greco, Ryan Doody, Melissa Fusco, Kevin Dorst, Kenny Easwaran, Brian Talbot, Bob Beddor, Zach Barnett, Abelard Podgorski, Al Hájek, Nic Southwood, Joshua Schechter, David Christensen, Jamie Dreier, and audiences at the National University of Singapore, the Australian National University, and Brown University for valuable feedback.

3. You objectively epistemically ought to believe truth and avoid error.

4. The norms of epistemic rationality must be explained in terms of trying to have true beliefs.

5. The epistemic value of a person's belief state is determined by its distance from the truth.

6. Rationality roughly consists in minimizing this distance to the extent possible, given one's limited information.

In previous work (Carr, 2017), I've raised some challenges for claim 6, and therefore for Accuracy-First Epistemology (AFE). The challenges go deeper, though. This paper argues against all of them. My target is a whole neighborhood of claims about epistemic norms associated with truth.

I focus on the example of claim 3, which I call the 'Truth Norm'. That's because the tools of deontic logic turn out to be convenient for explicating the problem. Proponents of many of these theses, e.g. accuracy-first epistemologists, don't generally talk in terms of the objective epistemic *ought* most of the time. But they still claim that epistemic rationality involves aiming to hold beliefs that best approximate some ideal—omniscience, or maximal accuracy. Once we try to clarify what ideal we're meant to approach—what truths our beliefs should be aiming to approximate—we encounter the problems I raise. Suppose you're more comfortable with the language of epistemic utility theory, for example. Then you can generate a translation schema between questions I raise about worlds where you do as you objectively epistemically ought and corresponding questions about the ideal, maximally accurate credence functions for one to adopt.

The contention of this paper is that the Truth Norm isn't a genuine epistemic norm. Why? Not merely because we ordinary believers can't satisfy it (in some sense of 'can't' that's sensitive to our cognitive limitations), or because it isn't action-guiding. Rather, the reason is that it's open to a variety of interpretations, and each interpretation, once made precise, yields unacceptable verdicts about what we ought to believe.

The plan for the paper is as follows: in section 1, I introduce a simple, schematic modal semantics for 'ought', and briefly discuss the ways in which it might imply 'can.' In sections 2–4, I discuss three forms of interpretation of the Truth Norms that we can elucidate within modal semantics: *world-relative*, *absolute*, and *mixed*. I show that each form of interpretation generates unacceptable predictions. In section 5, I consider a weaker variant of the Truth Norm—the Accuracy Norm—and show that it also makes unacceptable predictions. I also show that its problems generate worries for a widely accepted form of epistemic decision theory and other forms of epistemic consequentialism.

On this basis I conclude: there are no objective epistemic norms. In section 6, I discuss attempts to salvage some much weaker objective epistemic norms—norms that can't play the intended roles of the Truth Norm. I explain why even these should be rejected. There's no objective epistemic 'ought.'

1 A Truth Norm

1.1 The objective epistemic 'ought'

Philosophers often distinguish between subjective and objective 'ought's. The subjective 'ought' is sensitive to an agent's limited information; the objective 'ought' is not.[1] You *objectively* ought to take the very best option available to you, where what's best is insensitive to your beliefs or evidence. You *subjectively* ought to do what's rational in light of your beliefs or limited evidence. This involves aiming to approximate, in some sense, the recommendations of the objective 'ought.' For example, on a popular view, you subjectively ought to maximize expected utility. You objectively ought to maximize utility *tout court*.

A common thought in epistemology: the norms of epistemic rationality characterize the subjective *epistemic* 'ought'.[2] They provide the best means of approximating the objective epistemic 'ought'. The objective epistemic norm is widely thought to be the Truth Norm:[3]

> Truth Norm: you ought to believe all and only truths.

For short: you ought to be omniscient.[4] For different belief models, the Truth Norm is better stated in other ways. If we (plausibly) distinguish between rejecting a proposition and believing its negation, then the Truth Norm should be: believe all truths and reject all falsehoods. If our belief model includes credences instead of, or in addition to, full beliefs, then the Truth Norm should require having credence 1 in all truths and

[1] While I follow philosophers' convention of focusing on the modal 'ought', my discussion generalizes to other natural language normative expressions, e.g., 'should', 'have to', 'right', etc.

[2] N.B. I use 'epistemic "ought"' for the *deontic* necessity modals used in epistemology. The expression is unfortunately ambiguous: 'epistemic "ought"' is often also used to express (non-deontic) necessity given one's knowledge, e.g. 'With this warm front moving in, it ought to rain by tonight.'

[3] Some epistemologists might take the objective epistemic 'ought' to concern knowledge rather than true belief. The objections that I'll ultimately raise to the Truth Norm will apply equally to a knowledge norm, according to which one ought to know all truths (and, presumably, believe no falsehoods). For ease of exposition, I'll sometimes speak as though true belief is sufficient for knowledge.

[4] 'Omniscience', for our purposes, is stipulatively defined in terms of true belief, rather than knowledge.

0 in all falsehoods. I primarily discuss full belief, but my arguments are generalizable to all three models, and all three will be relevant in places.

The Truth Norm, on the intended interpretation, doesn't say that you should *become* omniscient. It says that you should already *be* omniscient. So the Truth Norm is best interpreted as time-coordinated: you ought$_t$ believe$_t$ all and only truths.[5] This coordinates the time when the required believing takes place with the time when the norm applies. If you ought, at t, to believe at t all and only truths, and at some later time, t^+, you come to believe all and only truths, you haven't thereby satisfied the obligations the Truth Norm imposed at t. You can only satisfy those by believing, at t, all and only truths. (Compare: while it's true that you ought$_t$ to pay$_{t^+}$ your parking tickets, it's also true that you ought$_t$ not have$_t$ parking tickets to begin with.)

1.2 'Ought' implies 'can'

A more typical objection to the Truth Norm comes from the charge that 'ought' implies 'can'. There's some sense in which we ordinary believers *can't* believe all and only truths.[6] Logical space includes uncountably many propositions: too many for cognitively limited believers like humans, with our little brains, to entertain. Moreover, some propositions might be impossible for cognitively limited believers to cognize. (For example, propositions that, in the believer's language of thought, cannot be mentally expressed within a lifetime.)

I'm not interested in this kind of objection. The objections to the Truth Norm that this paper raises, if correct, hold not only for ordinary believers but also for cognitively unlimited, ideal believers.

Even if we limit the set of propositions that the Truth Norm requires omniscience about, the problems I'll raise remain: we needn't assume that believers must have attitudes toward every proposition. Let W be the set of all possible worlds; any set of possible worlds is a proposition. Let $\mathcal{F} \subseteq \wp(W)$ be an agenda of first-order propositions that the agent can entertain. Assume that for all $p \in \mathcal{F}$, if a believer can entertain p, she can entertain Bp: the proposition that she herself believes that p). Let $\mathcal{A}$ be the smallest boolean closed set of propositions that includes $\mathcal{F}$ and also contains Bp for all $p \in \mathcal{F}$. The arguments below hold even if we assume that the agent can entertain only propositions in $\mathcal{A}$. So the believer might have only a small set of simple first- and second-order beliefs.[7] Throughout I assume that the believer has attitudes toward

[5] This doesn't time-relativize believed propositions, and the Truth Norm doesn't only apply to propositions about the present.

[6] See, e.g., Bykvist & Hattiangadi 2007.

[7] Possible worlds propositions introduce other controversies, but the big ones—logical omniscience, Frege's puzzle, closure under entailment—are not relevant to the discussion below.

some agenda of propositions $\mathcal{P} \supseteq \mathcal{A}$.[8] For simplicity, I assume that the agent has attitudes toward the same agenda of propositions in every possible world.

1.3 Deontic logic for the objective epistemic 'ought'

What does it mean to say that you ought to believe all and only truths? Standard deontic logic (SDL) provides an off-the-shelf, Kripke-style semantics for the objective epistemic 'ought':[9]

- Propositions are sets of possible worlds.
- $[\![\cdot]\!]$ is a valuation function assigning propositions to sentences. A sentence ϕ is true at w just in case $w \in [\![\phi]\!]$.
- A *deontic accessibility relation* R is a serial relation over possible worlds.[10]
- v is a *deontic alternative* to w iff wRv for the relevant deontic accessibility relation R. Deontic alternatives to w are ideal worlds with respect to relevant norms in w.[11]
- $[\![\text{Ought } \phi]\!]$ is true at w iff $[\![\phi]\!]$ is true at all deontic alternatives to w.

For ease of exposition, I'll move freely between object and metalanguage: for example, by re-expressing the previous sentence as ⌜*It ought to be that ϕ at w iff at all deontic alternatives to w, ϕ.*⌝

---

[8] $\mathcal{P}$ may contain contents of third+ order beliefs.

[9] For simplicity, I use SDL's naïve modal semantics. SDL faces a number of challenges for modeling natural language modal semantics, but these challenges mostly have to do with embeddings that are irrelevant for present purposes. The orthodox semantics for natural language deontic modals is Kratzer (1977, 1981). Kratzer's semantics derives the domain of quantification for deontic modals not from an accessibility relation, but from the combination of two contextual parameters: a *modal base*, determining a set of possibilities compatible with salient circumstances, and an *ordering source*, determining a preorder over worlds in terms of ideality relative to salient norms. The domain of quantification for the modal is roughly the set of highest ranked circumstantial possibilities.

One might worry the simplified deontic semantics I use neglects the modal base, and that any plausible circumstantial modal base for us humans will include the circumstance that omniscience is impossible. Alternatively, one might hold that the appropriate modal base is epistemic, and that any plausible epistemic modal base for us humans will include knowledge of our non-omniscience. If either of these hypotheses is correct, the Truth Norm is false for reasons independent of those I raise. But it's not clear that either of these constraints is defensible.

[10] R is *serial* iff $\forall w \in W(\exists v \in W(wRv))$.

[11] Different norms are therefore associated with different accessibility relations, representing, e.g., what's morally ideal, what's ideal for the purposes of winning our scavenger hunt, etc.

To interpret the Truth Norm, we need to specify the right accessibility relation for the objective epistemic 'ought' that it employs. In sections 2, 3, and 4, we'll consider a series of candidates.

2 World-relative interpretation

A first-pass interpretation of the Truth Norm: at any w, for any p in your agenda of propositions $\mathcal{P} \supseteq \mathcal{A}$, you objectively epistemically ought to believe p iff p is true at w. (Formally: where $\Box$ is the deontic epistemic necessity modal, for any $p \in \mathcal{P}$: $p \leftrightarrow \Box Bp$.) Call this the *world-relative interpretation* of the Truth Norm. It determines the following accessibility relation for 'ought':

> World-Relative: $\forall w, v \in W$, wRv iff for all $p \in \mathcal{P}$:
> $w \in p$ iff $v \in Bp$.

That is: v is a deontic alternative to w just in case in v you believe all and only truths about w.

The world-relative interpretation faces a problem. Because you're not omniscient, there's some true proposition p that you don't believe. So these two facts hold in any deontic alternative v to your world w:

(1) You believe p.

(2) You believe that you don't believe p.

The belief ascribed in (2) is a false belief at v, given that (1) is also true at v. So at all deontic alternatives to your world, you have a false belief. So you ought to have a false belief. But the Truth Norm says you ought to believe all and only truths. So the world-relative interpretation of the Truth Norm can't be correct.

Sorensen (1988) calls true conjunctions of the form $\ulcorner p$ *and I don't believe* $p \urcorner$ 'blindspots': they are possibly true propositions that are not truly believable (at least by agents who believe the conjuncts of conjunctions that they believe). Bykvist & Hattiangadi (2007) argue against the Truth Norm on these grounds: it violates an extremely weak 'ought'-implies-'can' principle, according to which you're obligated to ϕ only if it's metaphysically possible to satisfy your obligation.[12]

[12] The objection here is related to, but distinct from, both Moore's paradox (the impossibility of appropriately asserting: $\ulcorner p$, *but I don't know that* $p \urcorner$) and the Church-Fitch paradox (the impossibility of knowing some true propositions: $\ulcorner p$ *and no one knows that* $p \urcorner$).

Moore's paradox is generally taken to result from a knowledge norm on assertion. This makes no headway in addressing the worries I present for objective epistemic norms, because these worries aren't

This problem translates immediately to epistemic utility theory. Instead of looking to deontic alternatives to w, we look to ideal credence functions at w. Call a credence function Cr the "**accuracy touchstone**" for w just in case the accuracy of an agent's credences at w is measured in terms of its proximity to Cr. A world-relative accuracy touchstone for w is Cr^w_{WR}:

$$Cr^w_{WR}(p) := \begin{cases} 1, & \text{if } p \text{ is true at } w \\ 0, & \text{otherwise} \end{cases}$$

But since you're non-omniscient, there's some true p you don't assign credence 1 to. So for you, adopting Cr^w_{WR} means assigning credence 1 to p and assigning credence 1 to the proposition *I don't assign credence 1 to p.* These credences are necessarily self-falsifying. Insofar as the epistemic value of a credence function is accuracy, adopting these credences can hardly be alethically ideal. If the value that epistemic decision theories aim to maximize is accuracy in one's credences, that value can't be characterized in terms of proximity to self-falsifying credences. So the accuracy touchstone for a world shouldn't be Cr^w_{WR}.

---

ultimately about language.

The Church-Fitch paradox is more closely related to the objection above. Church (1945/2009) showed that on all normal modal logics, (ia) entails (ib), and the same form of argument shows that (iia) entails (iib).

(i) a. All truths are knowable.
b. All truths are known.

(ii) a. All truths are truly believable.
b. All truths are truly believed.

Still, there are key differences. The Church-Fitch paradox generates problems for (ia), because (ib) is false. Does the analogous entailment cause problems for the Truth Norm? Unlike (ib), the Truth Norm is a deontic claim. Even if (iib) is *actually* false, one might think it *ought* to be true. Indeed, this is one way of interpreting what the Truth Norm prescribes, which we'll explore in the next section.

The relation between the Church-Fitch paradox and the Truth Norm hinges on interpreting the latter as claiming: all truths are such that they ought to be truly believed (i.e., using the world-relative interpretation). If we combine this with a weak 'ought'-implies-'can' principle—⌜*ought* ϕ⌝ implies ⌜*logically possibly* ϕ⌝—then the Truth Norm can't apply to any non-omniscient agent.

But the Truth Norm is an informal claim, open to other interpretations. These interpretations hold some promise of better capturing the intuition behind the Truth Norm; none will be relevant for the Church-Fitch paradox.

3 Absolute interpretation

3.1 A different accessibility relation

The problem with the world-relative interpretation is that it assumes that in any ideal world v, you believe all and only truths about some other world: the nonideal reference world w. But if in v you believe all and only truths about w, then you are not omniscient with respect to v.[13]

The intuition behind the Truth Norm might be something like: ideally, you'd be omniscient. So any deontic alternative v to w should be one in which you are omniscient about all facts in v (rather than the facts in w). (Formally: for any $p \in \mathcal{P}$, $\Box(p \leftrightarrow Bp)$.) To capture this interpretation of the Truth Norm, we need a different accessibility relation. Call this the *absolute interpretation*:

Absolute: $\forall w, v \in W$, wRv iff $\forall p \in \mathcal{P}, v \in p$ iff $v \in Bp$

In English: the deontic alternatives to any w are the worlds v in which you're omniscient (i.e., in v, you believe all and only truths of v). This is absolute, rather than world-relative, because all possible worlds will share the same set of deontic alternatives.

On this accessibility relation, if v is accessible from w, then v is also accessible from itself. In other words, this accessibility relation satisfies *shift-reflexivity*. On all shift-reflexive frames, the following schema is valid: $\ulcorner\Box(\Box\phi \rightarrow \phi)\urcorner$. In English: it's obligatory that obligations be satisfied. In ideal worlds, one satisfies one's obligations; otherwise, ideal worlds would find themselves wanting. If even v finds itself wanting, then why would it be ideal to w? So w regards v as ideal only if v regards itself as ideal.

Notice that the world-relative interpretation doesn't satisfy shift-reflexivity. This was the heart of its problem: it entailed that an objectively epistemically nonideal world w will only regard another world v as ideal if you have false beliefs in v, and therefore v won't regard itself as ideal. So on that interpretation, the Truth Norm requires you to violate the Truth Norm.

The absolute interpretation nevertheless faces its own problems.

3.2 Problem #1

On the absolute interpretation, the deontic alternatives to your (nonideal) world are restricted to worlds in which you're omniscient. Your omniscience is a fact at all deontic alternatives, and so you believe that you're omniscient at all deontic alternatives.

[13] Assumption: there are no distinct possible worlds that make true exactly the same propositions.

So at your world, you objectively epistemically *ought* to believe that you're omniscient. But because this world is nonideal, the proposition that you're omniscient is not true. So on the absolute interpretation, you're required to believe something that's false.

Moreover, insofar as you know you suspend judgment on some propositions, on this interpretation, the Truth Norm requires you to believe something you *know for a fact* to be false, on the basis of introspection alone.

On the two interpretations we've considered, either you're required to have a false belief (on the world-relative interpretation) or you're required to believe something that's false (on the absolute interpretation). Formally:

- *World-relative:* $\exists p \in \mathcal{P}\Box(Bp \wedge \neg p)$
- *Absolute:* $\exists p \in \mathcal{P}(\Box Bp \wedge \neg p)$

Each of these conflicts with the intuitions motivating the Truth Norm.

This objection to the absolute interpretation isn't airtight, however. One might respond: the absolute interpretation predicts that you objectively ought to believe you're omniscient—but that's because you objectively ought to *be* omniscient. If you were doing what you objectively ought, then you would recognize your own omniscience. Perhaps we balk at the idea that you ought to believe that you're omniscient only because we mistakenly interpret the prescription subjectively. (Perhaps on the grounds that you're in a position to *know* that you objectively ought to believe you're omniscient.)[14]

3.3 Problem #2

A second, more serious problem arises for the absolute interpretation. Consider some contingently true proposition: for example, that there are cats. One might expect that, according to the Truth Norm, you objectively ought to believe that there are cats. But on this accessibility relation, this isn't the case!

Why? The absolute interpretation says that *any* world in which you're omniscient is a deontic alternative to the actual world. Even though [[there are cats]] is true at w, it isn't true at all deontic alternatives to w. (Presumably, cats are not a necessary condition on your omniscience.) So at some deontic alternatives to w, you don't believe that there are cats. So, on our modal semantics, it's not the case that you objectively epistemically ought to believe that there are cats.

[14] This defense of the absolute interpretation faces the worry that it implausibly severs any bridge between the subjective and objective 'ought'. The following bridge principle is plausible: if you know that you objectively ought to ϕ, and know that if you were to ϕ, you would do so knowingly, and know that ϕing weakly dominates all other options, then you subjectively ought to ϕ. The absolute interpretation entails that even this weak bridge principle is false: for example, when ϕ = *be omniscient*, or *truly believe you're omniscient*.

More generally: the set of worlds where you're omniscient is orthogonal to most ordinary, contingent, first-order propositions. And so on the absolute interpretation, the Truth Norm won't require you to believe these propositions. It only requires that you believe propositions entailed by your omniscience.

In short: the absolute interpretation isn't consistent with the Truth Norm, in letter or in spirit. The absolute interpretation predicts that there are some falsehoods that you ought to believe. And it predicts that, for most true contingent propositions p, it's not the case that you ought to believe p.

Again, these problems translate immediately to epistemic utility theory: the absolute accuracy touchstone assigns credence 1 to a proposition that's false at w (problem #1). It's also maximally imprecise over all propositions that aren't entailed by your omniscience, with the result that all these propositions make no contribution to your accuracy (problem #2). So this doesn't seem like the alethic ideal to approximate.

4 Mixed interpretations

4.1 Mixed interpretation #1

The second, more serious problem for the absolute interpretation resulted from the assumption that the deontic alternatives to w are *all* worlds where you're omniscient. That includes worlds where the contingent facts are very different from the actual world. So it includes worlds where your omniscience requires belief in many propositions that are false at the actual world. But what if we restrict the set of deontic alternatives to avoid this?

A first pass: perhaps the deontic alternatives to w are worlds where you're omniscient but the rest of the contingent facts are the same as at w.

> Mixed Interpretation #1: For all $w, v \in W$, wRv iff
>
> (i) for all $p \in \mathcal{P}$, $v \in p$ iff $v \in Bp$, *and*
>
> (ii) for all $p \in \mathcal{P}$ such that $p \not\models$ *you're not omniscient*, $v \in p$ iff $w \in p$.

I call this a 'mixed interpretation' because the accessibility relation restricts the deontic alternatives suggested by the *absolute* interpretation in a *world-relative* way.[15]

[15] In epistemic utility theory, this means that the mixed$_1$ accuracy touchstone for w is one which assigns credence 1 to all p s.t. $w \in p$ and $p \not\models$ *you're not omniscient*, and is transparent to itself (i.e., Cr^w_{M1}(your credence in p is n) $= 1$ iff $Cr^w_{M1}(p) = n$.)

Clause (i) ensures your omniscience in any deontic alternative v to w, while clause (ii) ensures that v and w are alike in all facts compatible with your omniscience: for example, there are cats in v iff there are cats in w. The proposition *there are cats* doesn't entail *you're not omniscient*, and is true at the actual world. So by clause (ii), it's true at all deontic alternatives to the actual world. So by clause (i), you truly believe it at all deontic alternatives to the actual world. And so you objectively ought to believe that there are cats.

Mixed interpretation #1 looks promising. It still predicts that you ought to believe you're omniscient—which may or may not be a problem (see section 3)—but it's able to predict that you ought to believe most true propositions at the actual world that aren't about your beliefs. So it apparently improves over both the world-relative and absolute interpretations.

But this interpretation still yields unacceptable verdicts. Consider (3):

(3) You're omniscient or Hillary Clinton is the US president.

(3) doesn't entail your non-omniscience, and (3) is not true at the actual world. So, by clause (ii), at any deontic alternative v to the actual world, (3) is false. But clause (i) ensures that any deontic alternative v to any world is a world where you're omniscient. So at every deontic alternative, (3) is true. So there are *no deontic alternatives* to the actual world.[16]

One might respond: if the actual world has no deontic alternatives, that can hardly be a problem for the Truth Norm. The absence of deontic alternatives entails the (trivial) truth of the Truth Norm.

But it also entails the trivial truth of the Falsehood Norm: believe all and only falsehoods. So the trivial truth of the Truth Norm should be no consolation to its proponents. More generally: it would be *metaphysically impossible* to do as one ought. Nothing would be permissible.[17]

4.2 Mixed interpretation #2

Mixed interpretation #1 is too strong, and has a natural weakening:

> Mixed Interpretation #2: For all $w, v \in W$, wRv iff
>
> (i) for all $p \in \mathcal{P}$, $v \in p$ iff $v \in Bp$, *and*

[16] For epistemic utility theory: w has no accuracy touchstone.

[17] This is why possible worlds semantics for SDL places a seriality condition on all deontic accessibility relations.

(ii) for all $p \in \mathcal{P}$ s.t. $p \nvDash$ *you're not omniscient*, if $w \in p$, then $v \in p$.

Mixed interpretation #1 uses a biconditional in its second clause, while mixed interpretation #2 uses a conditional.

Given this accessibility relation, from the fact that (3) is not true in the actual world, we can't infer that it's not true in the deontic alternatives to the actual world. So this interpretation avoids the problem we posed for its predecessor.

But this variant faces its own problems. Many worlds, including the actual world, will still have no deontic alternatives. To see this, let **n** be the sequence of actually winning Powerball numbers tomorrow. Here are three propositions true of the actual world:

(4) **n** will be the winning Powerball numbers tomorrow.

(5) If I believe that **n** will be the winning Powerball numbers tomorrow, then I'll buy a Powerball ticket today with **n**.

(6) I won't buy a Powerball ticket today.

None of these entails that I'm not omniscient; so all are true at any deontic alternative to the actual world. But they jointly entail that I fail the omniscience condition at any deontic alternative to the actual world. ((5) and (6) entail that I don't believe (4), but (4) is true. So I fail to believe something true.) But again, clause (i) requires any deontic alternative to be one in which I'm omniscient. So our second mixed interpretation also predicts that the actual world has no deontic alternatives.[18]

4.3 Mixed interpretation #3

We can avoid the problem for mixed interpretations #1 and #2 if we can guarantee there to be a nonempty set of deontic alternatives to any world.

The natural way to do this is to use a selection function. Let O be the set of worlds where you're omniscient. The selection function $\sigma : W \longrightarrow O$ maps each world to the uniquely closest (that is, most similar) world where you're omniscient.[19] This ensures that the deontic alternative to w is very much like w: if there are cats in w, then there are cats in w's deontic alternative.

[18] Or again: the actual world has no accuracy touchstone.

[19] This presupposes that there aren't multiple worlds tied for most similar. If we were to allow such ties, then that would generate further problems for this interpretation, beyond those I discuss below. For propositions about which tied worlds differ, this version of the Truth Norm can't tell you what to believe.

MIXED INTERPRETATION #3: $\forall w, v \in W$, wRv iff $\sigma(w) = v$.

Suppose that the correct semantics for counterfactuals is Stalnaker's (1968). Then Mixed Interpretation #3 can be summarized with a slogan: what you *ought* to believe is what you *would* believe, were you omniscient.

Will this do the job? Let v be the nearest possible world to the actual world in which you're omniscient. What might v be like? It depends on the similarity relation over worlds. And determining an appropriate similarity relation over worlds (for some purpose or other) is often a theoretical minefield.[20] We'll consider various options.

4.3.1 Future similarity, future omniscience

Suppose the closest world where you're omniscient, v, is one where you keep your extra knowledge to yourself. After all, the future of the world would be profoundly different if you shared all of your knowledge. Telling the world about how to cure cancer or to halt climate change would set the future off on a quite different course from actuality. So if the future of v is similar to that of the actual world, you won't share this information.

This means that v is a world where you're weirdly secretive: you hide most of your knowledge from other people. You don't tell your mother the winning Powerball numbers; you don't share the cure for cancer. So here are some other truths in v which, since you're omniscient in v, you believe of yourself in v:

(7) I'm weirdly secretive.

(8) I'm ungenerous with my mother.

(9) I allow people to needlessly die of cancer.

This mixed interpretation therefore predicts that in the actual world, you objectively ought to believe these propositions, even though they're not actually true of you. But that's absurd.[21]

An important upshot: there is *no bright line between first-order and higher-order beliefs*, or between propositions about your beliefs and other propositions. This means that weakening the TRUTH NORM to apply only to first-order beliefs is a nonstarter.

[20] See, e.g., Lewis 1979.

[21] All versions of mixed interpretation #3 face another objection: that in the deontic alternative to the actual world, (i) is true, hence believed, hence something you actually ought to believe of yourself:

(i) Most of my beliefs are entirely unsupported by the evidence that I've been exposed to.

And this problem will arise, in different ways, for other interpretations of the Truth Norm that we'll consider.

4.3.2 Present similarity, future omniscience

Let's try an alternative similarity relation. Maybe we should hold fixed your actual present willingness to share your beliefs, even if that means that the future of the actual world's deontic alternative will be quite different from the future of the actual world.[22] Then you needn't believe (7)–(9), because those propositions wouldn't be true in the nearest omniscience world.

This tweak generates its own problems. For then these propositions will be true of you:

(10) I'll tell my mother the winning Powerball numbers.

(11) I'll disseminate the cure for cancer.

I assume these are actually false. Because these are true at the actual world's deontic alternative v, you believe (10) and (11) in v. So on this similarity relation, our present interpretation of the Truth Norm predicts that in the actual world, you ought to believe (10) and (11) of yourself. But that's absurd.

4.3.3 Present similarity, future non-omniscience

In the previous two similarity relations, you elected not to act on your omniscience, or you acted on it. In both cases, this meant that there were some propositions that weren't immediately about your omniscience and that differed in truth value between the actual world and its deontic alternative. Your omniscience infected facts outside of your beliefs, facts that the Truth Norm shouldn't require or even apply to. The resulting variants of the Truth Norm entailed that you should believe some propositions that are actually false *and* not required by the Truth Norm to be true. Can we construct a similarity relation that quarantines the Truth Norm away from such propositions?

Let's pay attention to time-indexing: you ought$_t$ to ϕ just in case you would ϕ at t if you were omniscient at t. We might use a similarity relation where you're omniscient *only at* t, after which your body of beliefs contracts back to where it was before. You therefore have no time to act on your omniscience, affecting the future. For example, you have no chance of disseminating the cure for cancer. Your failure to do so is neither secretive nor ungenerous: by the time you consider whether to act, t is past and you

[22] Cf. Lewis 1979.

no longer remember the cure for cancer. Then we can predict that you're not required to believe either (9) or (11); similarly for the other examples.[23]

Again, this tweak generates its own problems. Now (12) is true of you in the actual world's deontic alternative:

(12) I'll lose almost all of my knowledge within the next minute.

So you believe (12) in the actual world's deontic alternative. So you objectively ought to believe (12) in the actual world, where (12) is false.

4.4 Lessons

The lesson to be learned from each of these attempted mixed interpretations: it's impossible to cleanly separate out the propositions that are related to your true beliefs from those that aren't. Propositions (7)–(12) aren't entailed by, or immediately about, your omniscience. But were you omniscient, your omniscience would make them true (on one or another plausible similarity relation), which would affect your beliefs about them. So our counterfactual interpretation of the Truth Norm predicts that you actually ought to believe these propositions, even though they're false. This can hardly capture the idea that you ought to believe all and only truths.

Notice that this objection differs in important respects from worries about the idea that you ought to believe:

(13) You're omniscient.

This is a falsehood (alas!), but as we suggested in section 3, requiring you to believe (13) might not be a problem. If the Truth Norm is correct, you ought to both be omniscient and recognize your omniscience. The only reason why (13) is false in the actual world is that you're not living up to your objective epistemic obligations.

But the same can't be said of (7)–(12). If the Truth Norm is correct, it doesn't require you to be weirdly secretive, or to tell your mother the winning Powerball numbers, or to lose almost all of your knowledge in the next minute. The fact that (7)–(12) are false is no epistemic shortcoming of yours. So on these interpretations, the Truth Norm requires you to believe false propositions that the norm doesn't require to be true.

Taking stock: we've canvassed several *prima facie* plausible interpretations of the Truth Norm; all have had unacceptable consequences. If you're with me so far, interpreting the Truth Norm should be starting to seem like a degenerating research program. On this basis, I tentatively conclude: there's no credible interpretation of the Truth Norm. And so there's no Truth Norm on belief.

[23] Thanks to [*omitted*] for suggesting this hypothesis.

5 Approaching truth

5.1 Accuracy maximization

Could there be a plausible alternative to the Truth Norm for the objective epistemic 'ought'? A thought: maybe the Truth Norm went wrong by demanding perfect omniscience. Instead, we might take the objective epistemic 'ought' to demand *maximizing* the accuracy of one's total doxastic state to the extent possible. Compare with decision theories like expected utility theory: there, the subjective 'ought' prescribes maximizing expected utility, and the objective 'ought' prescribes maximizing utility *tout court*.

So similarly, we might evaluate total doxastic states for degrees of accuracy. Where omniscience is *logically* impossible, one objectively ought to maximize the accuracy of one's belief state—even if in doing so, one's beliefs aren't perfectly accurate.

> Accuracy Norm: you ought to have the most accurate possible doxastic state.

The form of possibility should not be sensitive to limitations on your evidence (in contrast to the subjective 'ought').

It's unclear whether the resulting 'ought' can legitimately be thought of as objective, given that it constitutively involves agents' coping with uncertainty. But this 'ought' is, in some sense, more objective than the subjective epistemic 'ought', since it can require many beliefs that you don't already have, or even have evidential support for.

Nevertheless, worries loom. As with the Truth Norm, the Accuracy Norm has world-relative, absolute, and mixed interpretations. The Accuracy Norm pairs best with the *world-relative* interpretation. On that interpretation, we don't assume that deontic alternatives to the actual world are worlds in which you are alethically perfect, i.e., omniscient. On the absolute and mixed interpretations, deontic alternatives must be worlds in which you're omniscient, and so these interpretations face precisely the same objections as with the Truth Norm.

So we focus on the world-relative interpretation. The Accuracy Norm is meant to be distinct from, and to avoid the problems of, the world-relative Truth Norm. It can do so by requiring adoption of the maximally accurate beliefs at w that are compatible with being in w—where these might be less than omniscient about w.

> Accuracy Norm (World-Relative): In w, have the most accurate doxastic state compatible with being in w.

To see the problem with this proposal, consider a toy example. Suppose you only have doxastic attitudes towards these two propositions and their negations:

- p
- q =: I don't believe p.

We'll use a tripartite belief model that includes belief, suspension of judgment, and rejection. I assume that belief in truths and rejection of falsehoods are both more accurate than suspension of judgment on either truths or falsehoods, which is in turn more accurate than belief in falsehoods and rejection of truths.

Let w_{pq} be a world where both p and q are true. At w_{pq}, Accuracy requires you to adopt this belief state:

- believe q; reject *not-q*; and
- suspend judgment on p; reject *not-p*

Why is this the most accurate total belief state at w_{pq}, rather than believing all truths and rejecting all falsehoods at w_{pq}? Believing p is incompatible with being in w_{pq}: if you believe p, q is false, and hence you're not at w_{pq}. It's logically impossible to truly believe both propositions that are true at w_{pq}. Given that you can't have the perfectly accurate attitude toward p (belief), suspension is the best you can do; and then the rest of your attitudes can be perfectly accurate.

Problem #1. This maximally accurate total belief state is incoherent. If you reject *not-p*, coherence requires believing p. Yet in these circumstances, if you suspend judgment on both p and *not-p*, then your credences do not maximize accuracy: they are less accurate than the incoherent combination. Hence, the world-relative interpretation of the Accuracy Norm requires having incoherent doxastic attitudes.[24]

Problem #2. Your options for the maximally accurate beliefs to have at w_{pq} were limited by the fact that you could have attitudes toward a proposition about your belief state (q). So accuracy maximization requires holding fixed some facts about your own belief states at each world. This means that the more fully described worlds are within your space of propositions, the more constrained the options will be for maximizing accuracy. If you can entertain your doxastic attitude toward any proposition p, the maximally accurate doxastic attitude toward p at any world w will be *whatever doxastic attitude you happen to have* at w. If w is understood as a maximally specific possible world, determinate enough to specify your total belief state, then the most accurate belief state compatible with w is whatever your belief state at w happens to be. (This is, of course, the *only* belief state compatible with w.) So the relevant accessibility relation would be identity:

[24] See Caie 2013 and Castro & Vassend 2017 for a defense; Castro & Vassend use a credal analog of this example.

> Accuracy for Non-Toy Examples: $\forall w, v \in W$, wRv iff $v = w$.

It would therefore be *impossible* to violate the Accuracy Norm—rendering the norm trivial.

5.2 State-relative accuracy-maximization

An attempted remedy: we can distinguish *worlds* from coarser grained partitions of possible *states* of the world. Each state of the world is compatible with your being in multiple total doxastic states.

	s_1	s_2	…
adopt belief state 1	w_1^1	w_1^2	…
adopt belief state 2	w_2^1	w_2^2	…
⋮	⋮	⋮	⋱

Instead of a world-relative interpretation of the Accuracy Norm, we use a state-relative interpretation:

> Accuracy Norm (State-Relative): For any $s \in S$, if you're in s, you ought to have the most accurate doxastic state compatible with being in s.
>
> State-Relative: For all $w, v \in W$, wRv iff
>
> 1. $\exists s \in S$ s.t. $w, v \in s$, and
> 2. $\forall p \in \mathcal{P}$, if $s \subseteq p$, $v \in Bp$.

In English: v is a deontic alternative to w iff v is in the same state as w and in v, you're omniscient about all propositions entailed by that state.

Notice: whatever facts related to belief states are entailed by s will constrain what beliefs are maximally accurate at s. For example: if s entails both that p and that no one believes p, then the most accurate total belief state relative to s will demand suspension of judgment on p. So *prima facie*, this is no improvement over the world-relative interpretation. Because the partition of states can be coarser-grained, though, we can

require states to entail only first-order contents (i.e., no propositions about beliefs). So won't this solve the problem?

Problem #1: As section 4 showed, there's no clean way to separate out first-order vs. higher-order beliefs.

Problem #2: The only way to block the result that we always trivially have maximally accurate doxastic attitudes relative to our own world is to partition worlds into states that are silent about some of our doxastic attitudes. But this means that there are propositions that we are perfectly capable of entertaining, but maximizing accuracy doesn't require us to be correct about them.

What objectively ought we believe about our own beliefs? It's implausible that higher-order beliefs, and only higher-order beliefs, are *entirely unconstrained* by accuracy. So should we have states entail propositions about first-order beliefs? Second-order beliefs? Third? First-order beliefs about cats, but third-order beliefs about watermelons? This interpretation of the Accuracy Norm is partition-dependent, and any choice of partition will be arbitrary and unjustifiable. Different partitions will give incompatible verdicts.

5.3 Ramifications for epistemic decision theory

These problems for the Accuracy Norm, on either interpretation, infect any form of epistemic decision theory on which Perfectionism is true:

> Perfectionism: Our doxastic states are better to the extent that they approximate what we objectively ought to believe.

Perfectionism is widely assumed in epistemic decision theory.[25] If we objectively ought to have the maximally accurate doxastic state in a world or state, then:

Problem #1: We are sometimes, and perhaps always, objectively *and* subjectively required to be incoherent. Consider a credal variant of the example above: one has attitudes toward p and its negation, and toward q^* =: *No one has credence in p greater than .5,* and its negation. At w_{pq^*}, the maximally accurate credence function is c:

$$c(p) = .5; c(\neg p) = 0; c(q^*) = 1; c(\neg q^*) = 0$$

If we do better at w_{pq^*} by approximating c, then as Castro & Vassend (2017) show, c^* is *accuracy-dominated*:[26]

[25] I take the name from Pettigrew (2016); while his statement of the principle is slightly different, his (2018) makes clear that he accepts the thesis as I state it.

[26] I don't have the space to reconstruct Castro & Vassend's proof. Note that their argument is meant to refute Joyce's (1998; 2009) arguments for probabilism, and that the dispute may be reconstructed as a disagreement about which interpretation of the objective epistemic 'ought' to feed into Perfectionism.

$$c^*(p) = .5; c^*(\neg p) = .5; c^*(q^*) = 1; c^*(\neg q^*) = 0$$

Castro & Vassend conclude that c^* is always irrational, for any p, relative to any evidential situation. But this is an enormous bullet to bite: surely there are evidential situations in which you have very good reason to have credence .5 in both a proposition and its negation, and confident that everyone you're quantifying over has this credence. (Every ordinary toss of a fair coin, for example!) In short: not only will incoherence be objectively required, as in section 5.1, but it will also be subjectively, *rationally* required.

Problem #2: In some cases, a world or state will exhibit ties among distinct maximally accurate doxastic states. How do we determine which we should use to measure the inaccuracy of other doxastic states? These choices will have ramifications for what credences are rationally permissible. For example, suppose the agent has credences only toward p, $q^\dagger$, r, and their negations:

- p
- r
- $q^\dagger$ =: no one has credence greater than .5 in both p and r.

At $w_{pq^\dagger r}$, where the three propositions above are true, the following credence functions are tied for maximizing accuracy: c_p, which assigns credence .5 to p and perfectly accurate credences to everything else, and c_r, which assigns .5 to r and perfectly accurate credences to everything else. Which one should be treated as *the* ideal credence function at $w_{pq^\dagger r}$, relative to which the inaccuracy of other credence functions is measured? If we choose c_p, then for this space of propositions, it's always irrational to have credence .5 in both p and its negation; if c_r, then it's always irrational to have credence .5 in both r and its negation. Any decision between these will be *ad hoc*; and so the predictions of epistemic decision theory will be *ad hoc*.

6 There is no objective epistemic 'ought'

We've looked hard, but found no plausible way of understanding the Truth Norm. Each time we try to spell out an interpretation of the Truth Norm, we discover that it places absurd requirements on believers. Similarly for the Accuracy Norm. So I inductively conclude: there is *no truth or accuracy norm* on belief.

Now, the Truth and Accuracy Norms seem to be the weakest viable candidates for a general objective epistemic 'ought'. (For example, knowledge-based norms would

face the same problems.) If so, and if even these have counterexamples, then there is *no objective epistemic 'ought'*.[27]

One might think this is a wild overgeneralization: surely there are other, still weaker candidates for objective epistemic 'ought's. I'll briefly survey a few candidates.

6.1 Falsity avoidance

A first candidate:

Falsity Avoidance Norm: You ought not have false beliefs.[28]

Unlike the Truth Norm, this doesn't require you to believe anything: *ipso facto*, it doesn't require you to believe that you're omniscient, or that you'll let people needlessly die of cancer, or any of the other implausible prescriptions of the Truth Norm.

Even so, this norm still has some arguably implausible verdicts. For example: at every deontic alternative, you have no false beliefs; so at every deontic alternative you do not falsely believe that you have any false beliefs. So this norm entails that it's impermissible to believe that you have any false beliefs. If you're anything like me, you have extremely strong inductive evidence that in the actual world, you do have false beliefs, though you don't know which. Suppose you do have false beliefs. On this proposal, it's objectively impermissible to recognize your own fallibility; you're prohibited from believing a truth.

Again, it's an open question whether this is a problem, since it's only true that you have false beliefs because you violate the Falsity Avoidance Norm. Some may also view Preface Paradoxical belief states like these as subjectively impermissible, which takes the sting out of their objective impermissibility.

But more problematically: this norm is so weak that it can't play many of the roles that one might want an objective epistemic 'ought' to play: for example, what we ought subjectively to approximate in the pursuit of epistemic value, or what it would take to realize the greatest possible epistemic value. After all, the Falsity Avoidance Norm only requires avoiding false belief. You can satisfy this norm merely by suspending judgment on everything, believing no truths.

So suppose the Falsity Avoidance Norm is the strongest general objective epistemic norm. Then satisfying the Falsity Avoidance Norm is sufficient for doing as you objectively epistemically ought. This is inconsistent with two claims that many contemporary epistemologists accept—especially those who have been tempted by the Truth Norm:

[27] If so, the epistemic 'ought' of rationality is either not subjective, or subjective with no objective counterpart. The difference is probably terminological.

[28] Thanks to [*omitted*] for an equivalent suggestion.

(i) The objective 'ought' tracks value maximization: you objectively epistemically ought to maximize epistemic value.

(ii) True belief has more epistemic value than suspension of judgment.

For suppose both of these are true. Then if you suspend judgment on everything, you don't maximize epistemic value: having some true beliefs, without trading in any false beliefs, will have higher epistemic value. And so suspending judgment on everything is insufficient for doing as you objectively epistemically ought.

Suppose we bite these bullets. Notice that universal suspension—global skepticism—is guaranteed to perfectly satisfy Falsity Avoidance. Forming a belief about anything, by contrast, risks violating Falsity Avoidance. So a form of weak dominance argument shows that, on this norm, forming any beliefs whatsoever is *irrational*. You're subjectively epistemically required to be a skeptic!

6.2 Local objective epistemic norms

The various norms we've considered aimed for full generality over all propositions. So there is no fully general objective epistemic norm. But perhaps we can formulate *local* objective epistemic norms: for example: 'If there are cats, you ought to believe that there are cats'.

The problem is that *ought* is (arguably) agglomerative: if *ought* ϕ and *ought* ψ, then *ought* $(\phi \wedge \psi)$. And so each of these local norms would together generate a global norm: you ought to believe all and only truths. And this faces all the problems argued above.

In other words: we'd still need to provide an accessibility relation for this flavor of deontic modal. And this accessibility relation would determine a set of objectively epistemically ideal worlds. But which ones? Whatever we choose will deliver general verdicts. And for all of the options we've canvassed, these verdicts have been unacceptable. It might be possible to avoid the problem by using separate accessibility relations for different propositions, insisting that assertions change the contextually salient accessibility relation; but this looks hideously *ad hoc*.

In any case, we can still explain the intuition without appeal to objective epistemic norms. We can predict the truth of claims of the form ⌜*if p, you ought to believe p*⌝; it's just that the 'ought' they use is *subjective*. In the consequent of a conditional, the body of information that a subjective 'ought' is sensitive to may be stronger than the agent's information.[29] In particular, it will entail the proposition p expressed by the antecedent of the conditional.[30] Relative to this body of information, you ought to believe p—and

[29] Or the speaker's information, or the conversational participants' shared information, or the assessor's information, or…

[30] See Carr (2015, 680–2) for a defense.

because the relevant body of information may still be far from omniscient, this 'ought' is subjective.

6.3 Become omniscient

One might think that the problem lay in our starting assumption: that the Truth Norm evaluates your beliefs as they are now. Rather than saying that you ought to *be* omniscient, even though you aren't, perhaps the correct truth norm says that you ought to *become* omniscient. This will avoid many of the original Truth Norm's problems. After all, it doesn't require you right now to believe any falsehoods ('I will disseminate the cure for cancer'). And if you become omniscient, then the relevant propositions won't be falsehoods.

We can construct an accessibility relation for this on the basis of our most promising candidate above, Mixed Interpretation #3. On that interpretation, the Truth Norm claimed: what you ought$_t$ to believe$_t$ is what you would believe$_t$ if you were omniscient at t. On our new variant—call it the Become Omniscient Norm—what you ought$_t$ to believe$_{t^+}$ is what you would believe$_{t^+}$ if you were omniscient at t^+, where t^+ is, say, a second after t.

What does this requirement amount to? In the actual world, I'm sad to report, you will not be omniscient a second from now. Does this mean that you will violate your epistemic obligations? It's unclear. The Become Omniscient Norm imposes present obligations on your future attitudes: you're obligated now, at t, to become omniscient in the future, at t^+. So you're not currently violating your current obligations. Since you won't be omniscient at t^+, you will violate your current obligations. But you won't violate your obligations at t^+, because at t^+ the Become Omniscient Norm only imposes requirements on your behavior at a future moment, t^{++}. So if the Become Omniscient Norm is the strongest objective epistemic norm, then bizarrely, there is never any fact of the matter about whether you're doing as you ought!

7 Value without norms

If there's no Truth Norm on belief, does that mean that epistemic norms are entirely independent of truth-related considerations? Well, no. The falsity of the Truth Norm and its various weakenings doesn't give a reason to reject the idea that true belief is *valuable*. We can (partially) order worlds in terms of alethic value (value deriving from truth or accuracy). Worlds in which you are closer to omniscient are in important ways *epistemically better* than worlds where you have more uncertainty, which are in turn better than worlds where you have more mistaken beliefs. Alternatively—and to my mind, preferably—we can (partially) order doxastic states, construed as abstracta (say, credence functions, or 'belief functions' from propositions to B or $\neg B$) in terms

of alethic value. For example, if a credence function c's credence in p is closer to the truth of p than c^*'s for any $p \in \mathcal{P}$, then c has greater alethic value than c^*.

It's just that this value has no deontic import. There's no objective norm associated with this value. We can't say that you ought to maximize the value of your beliefs. For example: let p be a true proposition that you don't believe. At the actual world, then, the most alethically valuable doxastic state includes belief in both p and q =: *I don't believe p*. But we can't conclude on this basis that you *ought* to have those beliefs. If you did, you'd be guaranteed to have a false belief.

Nevertheless, the comparative alethic value of worlds, or of doxastic states, may still constrain the *subjective* 'ought' of epistemic rationality. It may, for example, be that you subjectively ought not adopt a doxastic state that is alethically dominated by another.[31] And so even if there is no truth norm on belief, the norms of epistemic rationality may still be governed by truth-based considerations. If so, though, the way in which they are so governed will differ greatly from pragmatic decision theories and consequentialist theories of pragmatic rationality in general.[32]

How, then, should we think about the relation between objective and subjective norms? Objective *pragmatic* norms are generally understood in terms of omniscience: very roughly, what you objectively ought to do is what you subjectively ought to do if you were omniscient. But the objective 'ought' doesn't extend from the pragmatic to the epistemic. While the pragmatic realm makes room for both subjective and objective 'oughts', the epistemic realm only has room for the former; and as I showed in section 5, this has serious implications for the possibility of importing pragmatic norms (e.g. decision theories) into epistemology. We can use omniscience to structure objective prescriptions, but we can't prescribe omniscience.

Is it true to say that you should believe the truth, then? No.
White (2009)

References

Bykvist, Krister & Anandi Hattiangadi (2007). "Does Thought Imply Ought?" *Analysis* 67(4), 277–285.

Caie, Michael (2013). "Rational Probabilistic Incoherence." *Philosophical Review* 122(4), 527–575.

Carr, Jennifer Rose (2015). "Subjective *Ought*." *Ergo* 2(27), 678–710.

[31] That is, for any doxastic states d and d^*, if d has greater alethic value than d^* at every possible world, then you subjectively ought not adopt d^*.

[32] See [*omitted*] for elaboration.

Carr, Jennifer Rose (2017). "Epistemic Utility Theory and the Aim of Belief." *Philosophy and Phenomenological Research* 95(3), 511–534.

Castro, Clinton & Olav Vassend (2017). "Ideal Counterpart Theorizing and the Accuracy Argument for Probabilism." *Analysis* 78(2), 207–216.

Church, Alonzo (1945/2009). "Referee Reports on Fitch's 'A Definition of Value'." In Joe Salerno (ed.), *New Essays on the Knowability Paradox*, 13–20. Oxford: Oxford University Press.

Joyce, James M. (1998). "A Nonpragmatic Vindication of Probabilism." *Philosophy of Science* 65(4), 575–603.

Joyce, James M. (2009). "Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief." In Franz Huber & Christoph Schmidt-Petri (eds.), *Degrees of Belief*, vol. 342, 263–297. Synthèse.

Kratzer, Angelika (1977). "What 'Must' and 'Can' Must and Can Mean." *Linguistics and Philosophy* 1(3), 337–355. doi:10.1007/BF00353453.

Kratzer, Angelika (1981). "The Notional Category of Modality." In H.-J. Eikmeyer & H. Rieser (eds.), *Words, Worlds, and Contexts: New Approaches in World Semantics*, Berlin: de Gruyter.

Lewis, David K. (1979). "Counterfactual Dependence and Time's Arrow." *Noûs* 13(4), 455–476.

Pettigrew, Richard (2016). *Accuracy and the Laws of Credence*. Oxford University Press UK.

Pettigrew, Richard (2018). "Making Things Right: the True Consequences of Decision Theory in Epistemology." In H. Kristoffer & Jeffrey Dunn Ahlstrom-Vij (ed.), *Epistemic Consequentialism*, 220–240. Oxford: Oxford University Press.

Sorensen, Roy A. (1988). *Blindspots*. Oxford University Press.

Stalnaker, Robert (1968). "A Theory of Conditionals." In Nicholas Rescher (ed.), *Studies in Logical Theory (American Philosophical Quarterly Monographs 2)*, 98–112. Oxford: Blackwell.

White, Roger (2009). "Evidential Symmetry and Mushy Credence." In *Oxford Studies in Epistemology*, Oxford University Press.